

## Grid Forum

Charlie Catlett, Argonne National Laboratory

Ian Foster, Argonne National Laboratory and University of Chicago

March 2000

### Abstract

The Grid Forum (GF) was established in 1999 as a community forum for discussion of Grid technology issues and as a means of coordinating efforts, promoting reuse and interoperability, and sharing results. GF holds regular meetings, maintains a web site ([www.gridforum.org](http://www.gridforum.org)), and operates working groups where ongoing technical discussions are facilitated with specific objectives outlined in working group charters. Participation in GF is from within the community of researchers and practitioners engaged in research, development, deployment, and support activities related to high-capability distributed software systems, or “grids.” The scope of the applications that motivate these activities is quite broad, including for example high performance processing applications, distributed collaborative environments, distributed data analysis, and remote instrument control. *A defining characteristic is a perceived need for services beyond those provided by today’s commodity Internet.* Grid Forum participation as of early 2000 includes over 150 individuals from universities, Federal laboratories, and companies organized into nine focused working groups. This paper outlines the history and motivation for Grid Forum as well as its current activities.

### 1. Grid Forum History and Motivation

Emerging network applications in such areas as high performance computing, information analysis, and distributed collaboration involve the coordinated use of multiple geographically distributed resources. The development of such applications can be significantly simplified if various “high-level” functions are present “in the network”. A large community of researchers and practitioners is now engaged in research, development, deployment, and operational support of a new class of infrastructure that provides these high-level functions. This “Grid” infrastructure provides caching, authentication, resource discovery, resource scheduling, and other services, with these new services being viewed by “Grid” applications in a manner analogous to how previous generation applications viewed services such as TCP sockets or protocols such as HTTP. Some of these issues are discussed in [1] and [2].

During the late 1980’s and early 1990’s many universities and Federal laboratories began to explore methods for large-scale distributed computing, adopting the term “Metacomputer” to refer to an integrated distributed computing

environment. By the mid-1990’s the World Wide Web and rapid expansion of the Internet, combined with early success in building “Metacomputers,” inspired many communities to think in terms of a more ubiquitous, transparent distributed computing environment that would be useful beyond simply high performance computing.

*This goal of a more general purpose, transparently integrated, persistent distributed computing infrastructure is referred to as a “Grid” in reference to the ubiquitous, transparent, persistent electrical power grid.*

Today there are dozens of major projects aimed at building “grids” in the US and elsewhere, and hundreds of researchers and practitioners who are building and/or supporting specific component technologies ranging from resource management to programming tools to security. In this context, the persistent grid infrastructure is viewed as a suite of services and protocols from which various types of “grids” might be constructed. These might be targeted to a specific set of users or a specific set of resources, but the variety of approaches will be quite broad. What is required is a rich suite of building blocks, appropriately interoperable, so that

multiple solution sets are possible depending upon the needs of a particular “grid” project. As important, building these “grids” should be possible without reinventing common services. This does not imply a single approach to each possible service but rather a selection of approaches that can interoperate.

Grid Forum began as an exploratory “Birds of a Feather” session at SC98 in November 1998, intended to explore the potential for an effort to coordinate among these various efforts. At the time, and today, though many technical coordinating bodies existed (IETF, OMG, W3C, etc.), there were no groups where there was a natural fit for this particular community. Several objectives were outlined, including:

- Discussion of common interests, problems, and solutions,
- Identification and definition of “best practices” and standards to promote code sharing and interoperability, and
- Promoting grid technology.

Based on these initial discussions, the Grid Forum was proposed as a regular gathering with only as much structure as is necessary to ensure that the group itself persists, an open membership based on participation by individuals rather than by organizations, and a focus on common infrastructure and ideas rather than particular projects. Further, many attributes of IETF were identified as worth emulating, particularly the spirit of “rough consensus and running code,” recognizing that standards cannot be forced and that working code in hand is better than two design documents in the bush. As stated in [3]:

*“[IETF] intends that working groups encourage and support open and fair participation and thorough consideration of technical alternatives. The procedures described below are designed with these goals in mind. Within those constraints, working groups are autonomous and each determines most of the details of its own operation with respect to session participation, reaching closure, etc. The core rule for operation is that acceptance or agreement is achieved via working group ‘rough consensus.’ ”*

Grid Forum will meet three times a year and has met twice since the November 1998 session. During these first two workshops an organizational structure was developed and nine working groups were formed. The third Grid Forum meeting,

“GF3,” is scheduled for March 22-24, 2000 in La Jolla, California. Working groups have developed charters based on the IETF working group guidelines [4], ensuring that the group’s focus and objectives are clear and measurable.

GF working groups are investigating a combination of research topics related to distributed systems, best practices for the design and interoperation of distributed systems, and the development of Applicability Statements (AS) for the implementation of grid software systems using current and emerging Internet Technical Specifications (TS) and other standards. In cases where there are no applicable technical specifications or standards and where no other bodies (IRTF, IETF, W3C, etc.) are engaged in the creation of these, GF may create a group to either investigate related research topics or to create a proposed TS.

A draft document outlining the relationship between GF and the Internet Standards Process [5] is being developed in the context of discussions with members of the Internet Architecture Board (IAB) and Internet Society (ISOC).

## 2. Grid Forum Structure and Objectives

The objectives of GF encompass three interrelated but distinct areas. The first relates directly to development of “production” Grid infrastructures and involves the development of Informational (e.g. recommended practices, deployment examples, etc.) and Applicability Statement (AS; e.g. implementation guidelines) documents that are necessary for building distributed systems. These documents will often be concerned with the application and/or integration of technologies developed within other groups such as IETF or W3C.

Second, GF intends to serve as a forum for discussion and exploration of longer-term issues. Typically these involve topics related to “middleware,” meaning protocols and services that operate primarily “below” the application but “above” the Internet Protocol Suite. In many cases, these middleware topics center around specific implementation issues, seeking to ensure interoperability of higher-level systems such as resource schedulers, authentication and authorization schemes, or data caching/storage services.

Finally, there are areas where Grid systems require that Technical Specifications (TS) be developed that address topics not covered by IETF or other standards-making bodies. In this respect, GF will operate in a fashion similar to IETF, focusing on shorter-term issues of engineering and standards making.

GF Working Groups are expected to have the stable long-term membership needed to promote the collaboration and teamwork required for exploration of research and implementation issues as well as development of implementation agreements that support interoperation. .

### 3. Grid Forum Process and Organization

GF was initially run by a group of volunteers including working group chairs, however it was recognized that some form of organization would be necessary to ensure that the group would persist beyond the initial enthusiasm. The group determined that a general Chair and a Grid Forum Steering Group (GFSG) would be necessary, with some form of election or appointment by an advisory group as a mechanism for determining individuals to serve in these roles. Charlie Catlett (Argonne National Laboratory) was appointed to be interim Chair and to organize the initial GFSG that also includes Ian Foster (Argonne National Laboratory and University of Chicago), William Johnston (LBL and NASA IPG), and Andrew Grimshaw (Univ. of Virginia).

During the first year of operation, the Chair and Steering Group are tasked with organizing meetings, assisting working group chairs in launching and running their groups, and developing processes and structures for governance.

### 4. Grid Forum Working Groups

Working Groups are the activity centers in the GF. A Working Group will typically be created to address a research, implementation, or operational area related to infrastructure necessary for building “grids.”

Nine working groups are operating within GF as of early 2000. Each has developed a charter that describes the intent of the group, its objectives, and a general timeline for reaching those objectives. The working groups are intended to persist over time, periodically updating their charters as goals

are reached and new requirements emerge. GF maintains a central website ([www.gridforum.org](http://www.gridforum.org)) and each working group also maintains a website where the charters, relevant background material, draft documents, and information of interest to the members are available. Group discussions primarily take place using a group email distribution list which is archived at the website.

A Working Group may be established at the initiative of an individual or group of individuals. Anyone interested in creating a GF Working Group must submit a charter for the proposed group to the GF Chair along with a list of proposed founding members. The charter will be reviewed for approval by the GFSG. Criteria for approving a working group as well as guidelines for developing a working group charter are taken from [6]. In general, working groups should have a clear definition of a relevant work area that is important to the Grid community, would not be coordinated without GF involvement, and does not directly duplicate the work of other groups (within GF or elsewhere). The proposed members should also have the appropriate expertise and understanding to undertake the work, and there must be a reasonable role for GF to play in the determination of the directions of the technology.

### 5. Summary of GF Working Groups

Each working group is summarized below. Additional information is available through the GF website ([www.gridforum.org](http://www.gridforum.org)).

#### 5.1. Security

The charter of the security working group is to investigate security issues relating to Grid environments. Initial focus areas include APIs and protocols (for example, investigating whether GSS-API and the Grid Security Infrastructure used by some Grid groups are appropriate for Grid systems and applications), interoperability of public key infrastructure elements; authorization issues; and documentation of security guidelines.

The Security working group chairs are Randy Butler (NCSA) and Marty Humphrey (University of Virginia).

#### 5.2. Scheduling and Resource Management

This group is concerned with various issues relating to scheduling and resource management of Grid resources. These include the development of a reference architecture that identifies key resource

management components; the definition of interface standards for resource management and scheduling functions; and the development of a standard for advanced reservations of Grid resources. In each area, draft documents are under development.

The Scheduling and Resource Management working group chairs are Bill Nitzberg (NASA NAS) and Jenny Schopf (Northwestern Univ.).

### 5.3. Remote Data Access

The Data Access working group is promoting the development of standard implementation practices for the construction of data grids: inherently distributed systems that tie together data and compute resources, supporting all aspects of data access and data manipulation. An end-to-end system provides support for:

- Data handling - ability to read data from a remote data resource for use within an application
- Remote processing - ability to filter or subset a data set before transmission over the network
- Publication - ability to add data sets to collections for use by other researchers
- Information discovery - ability to query across multiple information repositories to identify data sets of interest
- Analysis - ability to use data in scientific simulations, or for data mining, or for creation of new data collections.

Data handling systems provide the mechanisms needed to support distributed data access across heterogeneous data resources in computational grids and data grids. The requirements for a data handling system can be expressed as "transparencies" that hide grid implementation details. Examples include data set name transparency, data location transparency, access protocol transparency, latency hiding or data access optimization, and authentication transparency or single sign-on systems. The Data Access working group will seek common implementation practices for creating data handling systems.

The Remote Data Access working group chairs are Micah Beck (Univ. of Tennessee-Knoxville) and Reagan Moore (SDSC).

### 5.4. Grid Information Services

The goal of this working group is to identify requirements for, and facilitate the development of,

interoperable models and mechanisms for the information services necessary for grid-based computing. Specific issues include:

- The definition of meta-models to allow the creation of interchangeable schemas.
- The definition of formalisms to describe such schemas uniformly by the various working groups.
- The definition of a mechanism to access information that is stored in the schema.
- Educational outreach in order to provide the working groups with the necessary knowledge to use the models, the formalisms, and the mechanisms.
- Verification of the feasibility of the models, formalism, and mechanisms together with deployment teams.

The Grid Information Services working group chair is Gregor von Laszewski (Argonne National Laboratory).

### 5.5. Advanced Programming Models

The goal of this working group is to identify and investigate programming methodologies that support the effective development of algorithms that perform well in grid environments. The group has laid out the following near-term objectives:

- Characterize the design space for grid programming models,
- Identify a representative set of "Challenge Problems" facing grid programmers, and
- Identify existing or proposed tools, models, or other mechanisms aimed at these "Challenge Problems."

The Advanced Programming Models working group chairs are David Bader (Univ. of New Mexico) and Craig Lee (The Aerospace Corporation).

### 5.6. Grid Performance

This working group is concerned with the general problem of performance in grid environments. Its specific immediate goals are to develop a general format for interchange of performance data between tools, and to create a taxonomy of existing and proposed tools for the monitoring of Grid resources and services. These goals result in the following projects:

- Define a schema for data formats for performance monitoring. This would be a common interchange format that tools could use to interoperate.
- Taxonomy / classification of performance monitoring and analysis tools.
- Survey of existing tools classified by the above taxonomy.
- Recommendations on the aspects of grid applications, services and resources that should be monitored.

The Grid Performance working group chair is Valerie Taylor (Northwestern Univ.).

### 5.7. Grid User Services

Each group that is deploying grid services must also develop mechanisms for supporting users of those services. This working group brings together leaders from a variety of grid deployment efforts to coordinate and promote best practices for supporting grid users. In addition, as many grid users are either unaware of whose services they are using or are using services from a number of organizations, the GF user services working group is attempting to create a more integrated user services community. The group has three general objectives:

- Create a “clearinghouse” of documentation, training, software links, and other information pertinent to grid users, developers, and support staff,
- Identify “best of breed” support tools and methods that improve the quality and efficiency of grid user support, and
- Advocate for the end-user, facilitating and acting as a link between grid users and grid technology developers.

The Grid User Services working group chairs are Rita Williams (NASA NAS) and John Towns (NCSA).

### 5.8. Applications and Tools Requirements

The Grid Forum Applications and Tools Requirements Working Group was founded to provide critical information on the needs and requirements of Grid Application developers and end-users as well as tool developers to the developers of Grid infrastructure and services. The idea is to define the requirements of emerging Grid applications and tools for infrastructure developers, and to create a forum in which application developers can share information and express their

needs. These tasks are being undertaken via the documentation of applications and tools usage scenarios and requirements.

The Applications and Tools Requirements working group chairs are Fran Berman (UCSD), Robert Hood (NASA Ames), and Piyush Mehrotra (NASA Langley).

### 5.9. Account Management

This working group is concerned with issues relating to accounting in Grid environments. Early topics of discussion include standard formats for exchanging account information and the requirements of a cycle trading service. This discussion is a near-term test case for market-driven resource accounting and management in the context of supporting the pricing and exchange of grid resources given fluctuations in availability, realizable performance, application requirements (e.g. high throughput versus fast response), and other factors.

The Account Management chairs are Bill Thigpen (NASA NAS) and Tom Hacker (Univ. of Mich).

## 6. **Acknowledgements**

This paper describes the work of a broad community of Grid Forum participants who have contributed to the structure and progress of Grid Forum. In particular, the working group summaries have been harvested and adapted from the work of the working group chairs.

## 7. **REFERENCES**

- 
- [1] Foster, I. and Kesselman, C. (Eds), “The Grid: Blueprint for a Future Computing Infrastructure,” Morgan-Kaufmann, 1999.
  - [2] Aiken, A., Strassner, J., Carpenter, B., Foster, I., Lynch, C., Mambretti, J., Moore, R., and Teitelbaum, B., “Network Policy and Services: A Report of a Workshop on Middleware,” RFC 2768, Feb 2000.
  - [3] Bradner, S., “IETF Working Group Guidelines and Procedures”, RFC 2418, September 1998.
  - [4] Ibid.
  - [5] Bradner, S., “The Internet Standards Process – Revision 3”, RFC 2026, October 1996.
  - [6] Ibid.